PUB: Product Recommendation with Users' Buying Intents on Microblogs

Xiaoxuan Ren[1,2](✉), Tianshu Lyu[1,2], and Yan Zhang[1,2]

[1] Department of Machine Intelligence, Peking University, Beijing 100871, China
{renxiaoxuan,lyutianshu,zhy.cis}@pku.edu.cn
[2] Key Laboratory of Machine Perception (MOE), Beijing 100871, China

Abstract. Recommendation systems mostly rely on users' purchase records. However, they may suffer problems like "cold-start" because of the lack of users' profiles and products' demographic information. In this paper, we develop a method called PUB, which detects users' buying intents from their own tweets, considers their needs, and extracts their demographic information from their public profiles. We then recommend products for users by constructing a heterogeneous information network including users, products, and attributes of both. In particular, we consider users' shopping psychology, and recommend products that better meet their needs. We conduct extensive experiments on both direct intent recommendation and additional product recommendation. We also figure out users' potential preference which can help to recommend a great varied types of products.

Keywords: Product recommendation · Buying intent detection
Heterogeneous information network embedding

1 Introduction

Most existing product recommendation systems use users' historical transaction records or interactions like rating or clicking, and mostly rely on techniques like collaborative filtering [3,9,15]. However, the mere collaborative filtering technique may suffer from the "cold-start" problem for the lack of information of a new user. Furthermore, collaborative filtering cannot precisely discover users' real needs either, especially the ad hoc ones. Actually a lot of information on OSN platforms (e.g. the tweets) contain users' buying intents [11], which can help to fill up the entrenchments.

In our work, we make a bridge between the e-commerce websites and social networks and develop a method, PUB, Product recommendation with Users' Buying intent. Figure 1 is the framework of PUB. PUB identifies users' buying intents, needs for the products and product adopters by a bootstrap based method from Sina Weibo[1], the Chinese largest OSN platform. Then PUB selects

[1] http://weibo.com.

H. Hacid et al. (Eds.): WISE 2018, LNCS 11233, pp. 308–318, 2018.
https://doi.org/10.1007/978-3-030-02922-7_21

the most appropriate product for the specific user from JD.com[2], one of the largest B2C e-commerce websites in China. We also develop a heterogeneous information network with nodes of users, products and attributes. Finally, PUB recommends products for users which fit their needs and also some other kinds of products that the users may like.

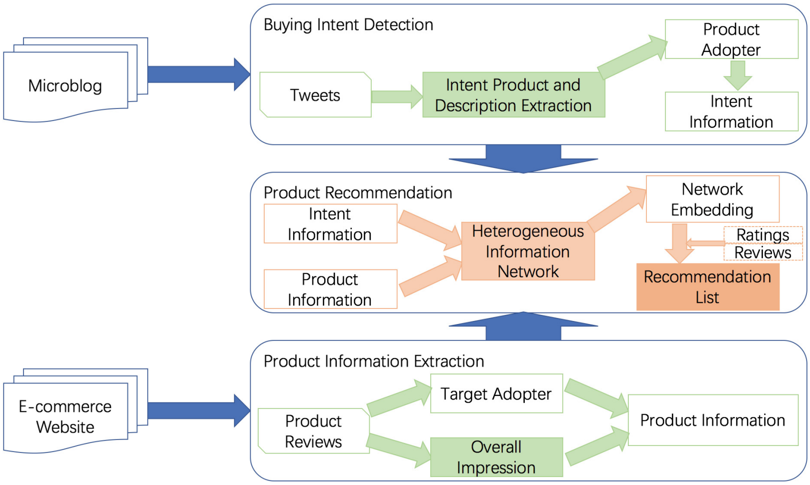

Fig. 1. Framework of PUB.

The main contributions of our method are as follows: (1) We accurately detect buying intents from users' real-time tweets, and extract users' needs and demographic information. (2) We develop a heterogeneous information network to link together e-commerce websites and social networks, and also consider users' shopping psychology to meet the users' demand. (3) We conduct extensive experiments on both direct product recommendation and additional product recommendation, and the results show that our method outperforms other baselines on extracting users' intents as well as recommending products.

2 Related Work

Social Network Extraction. Kröll *et al.* introduce the idea of intent analysis, and evaluate intent profiles from speeches [5]. Hollerit *et al.* detect commercial intent in tweets [4], and their work is considered a first step to bring together buyers and sellers. The most related work is from Wang *et al.* [11], in which they identify users' intents and classify them into six categories. However, they fail to extract users' needs and don't recommend products for the users at last. Meanwhile, there exist some works on extracting users' demographic characteristics.

[2] http://jd.com.

Bachrach *et al.* study how users' public profiles on Facebook relate to their personalities [1]. There are also studies on extracting users' public profiles in Weibo and using these information for product recommendation [17], but they fail to extract the potential preference of users.

Product Recommendation. Most recommendation systems rely on collaborative filtering [3,9], which suffer the "cold-start" problem, and mainly takes into account ratings of users on products and users with similar interests. Xu *et al.* propose a semantic path based personalized recommendation to predict the rating scores of users on items [13]. Zhao *et al.* develop a novel recommendation system based on matching the users' demographic information with product demographics [17]. Some researchers have begun to notice the heterogeneous information for recommendation. Shi *et al.* propose a weighted heterogeneous information network with meta path called SemRec to predict the rating scores of users on items [10].

3 Preliminary Concepts

In this Section, we describe some preliminary knowledge as well as the notations used in this paper.

Buying Intent Detection. According to the description of a commercial intent tweet [4], we define a tweet with buying intent if it (1) contains at least one verb and (2) explicitly describes the intention to buy a certain product (3) in a recognizable way. Here is an example:

Please recommend a bright light lipstick, my girlfriend wants to buy one.

The tweet above contains buying intent which satisfies all the three conditions in the definition. However, we can figure out that the **product adopter** is not the user but the user's girlfriend. Inspired by the definition on intent-indicator and intent-keyword [11], we define intent-indicator, intent-product and intent-description as below. **Intent-indicator** is a verb phrase or infinitive phrase that express the users' intent to buy something. "Lipstick" in the example is an **intent-product**, which is a noun in most cases, and is the product that the user wants to buy. **Intent-description** is an adjective or noun between the intent-indicator and intent-product, or appearing alone.

User and Product Demographics. User demographics describe the intent buyers' characteristics. In our work, we only consider the attributes of age and gender, for users are not willing to provide other attributes in most cases. The gender is classified into male and female, while the age is grouped into eight clusters: 0–3, 4–6, 7–17, 18–29, 30–40, 41–55, 56–64, 65+. The product demographics are extracted from product reviews, which describe the characteristics of the buyers of the product.

Heterogeneous Information Network. Heterogeneous information network have been proposed as a general data representation for many different types of data [16]. A heterogeneous information network has different types of objects and different links representing different relations, which can better fit in most scenarios, especially in product recommendation. Similar to [14], we define heterogeneous information network as follows:

Let $G = (V, E)$ *denotes a graph with an object type mapping function* $\phi : V \rightarrow A$ *and a link type mapping function* $\psi : E \rightarrow R$*(*$|A| > 1$ *or* $|R| > 1$*). Each object* $v \in V$ *belongs to one object type* $\phi(v) \in A$*, and each link* $e \in E$ *belongs to a relation type* $\psi(e) \in R$*.*

4 Methodology

4.1 Buying Intent Detection and Users' Demographic Information Extraction

Let u and $U = \{u_t\}_{t=1}^{M}$ denote a user and the whole user set. Similarly, let p and $P = \{p_t\}_{t=1}^{N}$ denote a product and the whole product set. To detect $u's$ buying intents on Weibo, we first filter out irrelevant tweets with a list of key words, then we use a bootstrap based method proposed in Algorithm 1.

First of all, we input a microblog tweet sentence corpus $\mathcal{T}$, and a seed set of intent-indicators, such as "want to buy". The output is an extension of the intent-indicators set $\mathcal{I}$ and intent-products set $\mathcal{P}$, as well as the set of intent-descriptions $\mathcal{D}$. The function **ExtractIntentProduct**(i, t) and **ExtractIntentDescription**(i, t) aim to extract users' intent-products and intent-descriptions according to the intent-indicators. We consider the positions of the intent-products may be the preceding n_1 tokens of the intent-indicators, the middle n_2 tokens or the following n_3 tokens. In order to better obtain users' preference, we also extract the intent-descriptions, which often appear before the intent-products or alone.

With the intent-products that frequently co-occur with the intent-indicators, we in turn use these intent-products to find more intent-indicators in function **ExtractIntentIndicator**(i, t). We repeat the previous steps, until no more intent-indicators or intent-products can be found.

The product adopters is extracted similar with buying intents extraction. Then we extract users' demographic characteristics from their public profiles.

The product adopter, needs for the product and $u's$ demographic characteristics are all regarded as the $u's$ attributes, denote as a_u and $A_u = \{a_t^u\}_{t=1}^{L}$.

4.2 Products Demographic Information Extraction

We use the online products reviews from JD.com to extract products demographic information. First of all, we find the target adopter of product p through the bootstrap based method mentioned in Sect. 4.1, and infer the age and gender. After that, we learn the buyers' overall impression of the products. To achieve it, all the reviews of product p are merged into a single document. We segment Chinese streams into words and extract the keywords from the reviews of one product with TextRank [6].

Algorithm 1. Bootstrap based algorithm for buying intent detection.

Input: microblog tweet sentence corpus $\mathcal{T}$; seed intent-indicator patterns
Output: an extension set of intent-indicator patterns $\mathcal{I}$; a set of intent-products $\mathcal{P}$; a set of intent-descriptions $\mathcal{D}$

$\mathcal{I} \leftarrow$ seed intent-indicator patterns
$\mathcal{I}^{'} \leftarrow$ seed intent-indicator patterns
$\mathcal{P} \leftarrow \varnothing$
$\mathcal{D} \leftarrow \varnothing$
repeat
    $\mathcal{P}^{'} \leftarrow \varnothing$
    $\mathcal{D}^{'} \leftarrow \varnothing$
    for each pattern $i \in \mathcal{I}^{'}$ **do**
        for each sentence $t \in \mathcal{T}$ **do**
            if i exists in t **then**
                $\mathcal{P}^{'} \leftarrow \mathcal{P}^{'} \cup$ ExtractIntentProduct(i, t);
                $\mathcal{D}^{'} \leftarrow \mathcal{D}^{'} \cup$ ExtractIntentDescription(i, t);
            end if
        end for
    end for
    $\mathcal{I}^{'} \leftarrow \varnothing$
    for each product $p \in \mathcal{P}^{'}$ **do**
        for each sentence $t \in \mathcal{T}$ **do**
            $\mathcal{I}^{'} \leftarrow \mathcal{I}^{'} \cup$ ExtractIntentIndicator(i, t);
        end for
    end for
    $\mathcal{I}^{'} \leftarrow$ ExtractTopFrenquentIndicator$(\mathcal{I}^{'})$;
    $\mathcal{I} \leftarrow \mathcal{I} \cup \mathcal{I}^{'}$
    $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{P}^{'}$
    $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}^{'}$
until no new pattern is identified;
return an extension set of intent-indicator patterns $\mathcal{I}$;a set of intent-products $\mathcal{P}$;a set of intent-descriptions $\mathcal{D}$

Let $G = (V, E)$ be a graph with the set of vertices V and set of edges E. Let $In(V_i)$ be the set of vertices that point to vertex V_i, and $Out(V_i)$ be the set of vertices that V_i points to. Taking into account the edge weights, the score associated with a vertex in the graph can be defined as:

$$WS(V_i) = (1 - d) + d \sum_{V_j \in In(V_i)} \frac{w_{ji}}{\sum_{V_k \in Out(V_j)} w_{jk}} WS(V_j) \tag{1}$$

Where d is a damping factor that can be set between 0 and 1, w_{ij} is the weight between two edges.

We filter the words with Part-of-Speech (POS) including noun, adjective and descriptive. We also add some stop words. Let the words be the vertices of the graph, and build the edges with $co - occurrence$ relation: if two vertices' corresponding lexical units co-occur within a window of maximum N words, where N can be set anywhere from 2 to 10 words, then the two vertices are connected. We calculate the score for each word, and collapse the sequences of adjacent keywords into a multi-word keyword. At last, we export the top-10 words as the impression keywords.

The impression keywords, target adopter with age and gender are seen as the product $p's$ attributes, denote as a_p and $A_p = \{a_t^p\}_{t=1}^R$.

4.3 Graph Embedding for Product Recommendation

We create a heterogeneous information network consists of users, products, and their attributes. An illustration of our product recommendation in heterogeneous information network is shown in Fig. 2. What needs to be pointed out is that the attributes of users and products often overlap, then we let $A = A_u \cup A_p$ to denote the whole attribute set.

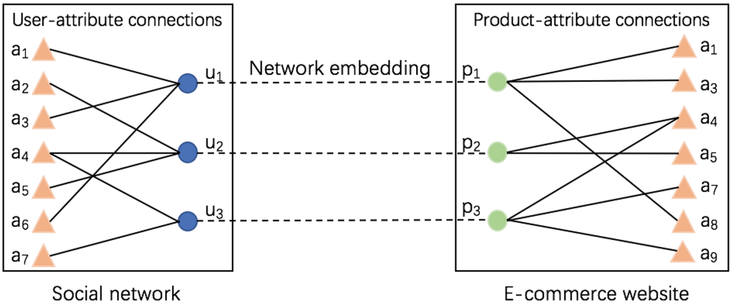

Fig. 2. A heterogeneous information network.

In our work, we use Deepwalk [7], which learns a latent representation of adjacency matrices using deep learning techniques developed for language learning. In order to better satisfy the users' needs, we also merge the words with similar meaning, like "dark" and "black". At last, we predict which products should be recommended to the user by calculating the cosine similarity between the given user node and the products. Then the input and output of the network are defined as:

Input: A social network domain with $\{U, A_u\}$; an e-commerce domain with $\{P, A_p\}$; and $A = A_u \cup A_p$.

Output: A personalized recommendation list for each user u with products by ranking the similarity.

We also consider the ratings and the numbers of reviews, because according to a new study from Psychological Science [8], people prefer more-reviewed items (even with poor ratings). We calculate the users' preference for a given product and denote it as Pr. The equation for calculating the preferred value between the user and the product is as follows:

$$Pr = Sim \cdot (\alpha \log_m(\#CM_p + 1) + \beta R_p) \tag{2}$$

Where $\#CM_p$ is the number of the reviews of the product p, R_p is the rating for the product (range from 0 to 5), α and β are weights of every term. The log base m can be calculated as below:

$$m = \mathbb{E}(\#CM)^{\frac{1}{\mathbb{E}(R)}} \tag{3}$$

where $\mathbb{E}(\#CM)$ is the expectation of the number of the reviews of the products, and $\mathbb{E}(R)$ is the expectation of ratings.

With a heterogeneous information network, we can not only recommend direct products that the users want (with the category label that the users' intents belong to), but also recommend additional products with all categories.

5 Experiment

5.1 Experimental Setup

We crawl tweets from Weibo which contains at least one keyword expressing buying intent, like "want to buy" or "please recommend", and get 6,095 tweets ultimately. We also crawl the users' demographic information of age and gender. We use the dataset from [12] which has 246,444 products and 12,127,267 users with 138,905,740 records of reviews.

5.2 Evaluation on Buying Intent Detection

Because there is no existing ground truth, we invite two annotators who are familiar with the language habits on microblogs, and ask them to label whether the tweets contain buying intent with the description in Sect. 3. If the two annotators have different ideas on one tweet, then we abandon this tweet. At last, we obtain 1,090 tweets with buying intents, and 4,867 tweets without buying intents. The Cohen's Kappa agreement coefficient between annotators is 93.5%, which is really satisfied.

We consider the following comparison methods in our experiments for the buying intent detection:

Wang's: Proposed in [11], which only consider the intent-products just following the intent-indicators with a bootstrap based method, and fail to consider the intent-descriptions.

SVM: We train the Support Vector Machine (SVM) with linear kernel, and use the bag-of-words of the tweets as the input characteristics for classifier.

Ours: Our proposed method for buying intent detection, which considers the intent-indicators, different possible positions for intent-products, and can extract users' needs for the products at the same time.

The results are shown in Table 1. From the results we can tell that our method outperforms the two baseline methods in terms of precision, recall and F_1 Measure. In particular, the improvement of recall is significant compared with Wang's method by 19%. We may conclude that considering the positions of intent-products and the existence of intent-indicators, we can extract more tweets with buying intents.

Table 1. Performance comparison on buying intent detection (%).

Methods	Prec.	Rec.	F_1
Wang's	74.07	27.03	39.60
SVM	70.45	41.89	52.54
Ours	82.93	45.95	59.14

5.3 Evaluation on Product Recommendation

We choose three frequently mentioned intent product types in the Weibo dataset, camera, lipstick, and foundation (makeup). The statistics of the three products are summarized below in Table 2.

Table 2. The statistics of the three given products.

Types	#brands	#products	#comments	#intents
Carema	41	2,341	623,081	30
Lipstick	89	476	432,326	74
Foundation	126	606	258,096	40

Then we find the target adopter, and learn the overall impression every product. Here is an example of the overall impression we learn from the reviews of a Bobbi Brown lipstick shown in Table 3.

Table 3. The overall impression of a lipstick.

Hydrated, wife, orange, light, warm colour, good quality, reasonable price, anthentic, pleasing colour, matte

However, the overall impression from JD.com of the same product is only one word "not bad", even with 62 reviews. We may infer that our learned overall impression are more suitable for the recommendation system, as it better conclude the image of the given product.

We invite five annotators to evaluate which ways recommended products better fit the intent needs. We give each annotator the needs for the product, the user's age and gender, the product's demographic information, the product's overall impression and the rating as well as the number of reviews. We also break the order of the three methods for the three products.

We consider three methods for product recommendation in our experiments:

MART [2]**:** A pointwise learning-to-rank approach. In this approach we only consider the profiles of age and gender extracted based on a given user and a candidate product.

PUB$_{Sim}$: This is a graph embedding method but only consider the similarity between the given user node and the products.

PUB$_{All}$: This is the proposed method which not only consider the similarity, but the ratings and the numbers of reviews as well.

Table 4. Performance comparison for chosen the most suitable product for the given intent from five annotators.

Annotators	Camera			Lipstick			Foundation		
	MART	PUB$_{Sim}$	PUB$_{All}$	MART	PUB$_{Sim}$	PUB$_{All}$	MART	PUB$_{Sim}$	PUB$_{All}$
A	6	9	15	21	22	31	9	10	21
B	10	7	13	16	22	36	11	10	19
C	10	8	12	19	27	28	13	12	15
D	6	8	16	25	18	31	6	11	23
E	12	10	8	18	24	32	13	10	17
Avg	8.8	8.4	12.8	19.8	22.6	31.6	10.4	10.6	19

From the results in Table 4 we can see that the products recommended from our final method outperforms others. We may conclude that the using of graph embedding can extract the users' latent preference. However, the learning to rank method only consider the relation between the query (buying intent tweet) and the relevant document (an adopted product).

5.4 A Case Study on Direct Product Recommendation and Additional Product Recommendation

Direct Product Recommendation. Firstly, we give an example on direct product recommendation, which means recommend the kind of products that the users want. We use the buying intent described in Sect. 5.2, that a girl wants to buy an orange lipstick, with the age between 18–29 and the gender as female. The recommended products' demographic information of the three methods are shown in Table 5.

Table 5. The products' demographic information recommended by three methods.

Information	MART	PUB$_{Sim}$	PUB$_{All}$
Age	18–29	18–29	18–29
Gender	Female	Female	Female
#reviews	39	15	784
Rating$_{avg.}$	4.74	4.87	4.65
Impression	Hydrated, girlfriend, cheap, light	Hydrated, orange, light, good quality, matte	Hydrated, orange, cheap, light, pleasing color, authentic

The product recommended by MART only satisfies the age and gender, while the products recommended by two proposed methods seems to be similar, but the $Ours_{All}$ method's product have more reviews and lowest rating.

Additional Product Recommendation. After the experiments mentioned above, we then build a heterogeneous information network with the whole 1,095 buying intents and 246,444 products, and try to figure out what will recommend to the user with a settled intent.

Table 6. The other product recommended for the girl demand for orange lipstick.

Information	Bag	Ear phone	Luggage
Age	18–29	18–29	18–29
Gender	Female	Female	Female
#reviews	934	592	845
$\text{Rating}_{avg.}$	4.28	4.57	4.32
Impression	Cheap, orange, able to hold things, portable	Fashion, good-looking, orange, good sounds	Durable, size suitable, light, orange

In Table 6, we can see that the method recommends some other products for the girl which are suitable for young ladies, and satisfy her own preference. This is quite useful in the real world, for it extracts the latent rations between different types of products, recommends the products that others in the same conditions may like, and people may want to buy a bag which meet her needs even though what she really wants at that time is a lipstick.

6 Conclusion and Future Work

In this paper, we develop a method called PUB, which can detect users' buying intents and profiles from microblogs, and recommend products from e-commerce websites. We use a graph embedding method, and construct a heterogeneous information network. We conduct experiments on both buying intent detection and product recommendation. We do the recommendation job on both direct product and additional product recommendation, to figure out the potential preference. As a future work, we want to incorporate more features considered into our method, especially those implicit features. We also plan to add sentiment analysis for the product reviews or users' tweets. The graph embedding method can be improved and optimized as well.

Acknowledgment. This work is supported by 973 Program under Grant No. 2014CB340405, NSFC under Grant No. 61532001, and MOE-ChinaMobile under Grant No. MCM20170503.

References

1. Bachrach, Y., Kosinski, M., Graepel, T., Kohli, P., Stillwell, D.: Personality and patterns of facebook usage. In: Proceedings of ACM Web Science (2012)
2. Friedman, J.H.: Greedy function approximation: a gradient boosting machine. Ann. Stat. **29**, 1189–1232 (2001)
3. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.S.: Neural collaborative filtering. In: Proceedings of WWW (2017)
4. Hollerit, B., Kröll, M., Strohmaier, M.: Towards linking buyers and sellers: detecting commercial intent on Twitter. In: Proceedings of WWW (2013)
5. Kröll, M., Strohmaier, M.: Analyzing human intentions in natural language text. In: Proceedings of K-CAP 2009 (2009)
6. Mihalcea, R., Tarau, P.: TextRank: bringing order into text. In: EMNLP (2004)
7. Perozzi, B., Al-Rfou, R., Skiena, S.: DeepWalk: online learning of social representations. In: Proceedings of ACM SIGKDD (2014)
8. Powell, D., Yu, J., DeWolf, M., Holyoak, K.J.: The love of large numbers: a popularity bias in consumer choice. Psychol. Sci. **28**, 1432–1442 (2017)
9. Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Item-based collaborative filtering recommendation algorithms. In: Proceedings of WWW (2001)
10. Shi, C., Zhang, Z., Luo, P., Yu, P.S., Yue, Y., Wu, B.: Semantic path based personalized recommendation on weighted heterogeneous information networks. In: Proceedings of ACM CIKM (2015)
11. Wang, J., Cong, G., Zhao, W.X., Li, X.: Mining user intents in Twitter: a semi-supervised approach to inferring intent categories for tweets. In: AAAI (2015)
12. Wang, J., Zhao, W.X., He, Y., Li, X.: Leveraging product adopter information from online reviews for product recommendation. In: ICWSM (2015)
13. Xu, G., Gu, Y., Dolog, P., Zhang, Y., Kitsuregawa, M.: SemRec: a semantic enhancement framework for tag based recommendation. In: AAAI (2011)
14. Yu, X., et al.: Personalized entity recommendation: A heterogeneous information network approach. In: Proceedings of WSDM 2014 (2014)
15. Zhang, H., Shen, F., Liu, W., He, X., Luan, H., Chua, T.S.: Discrete collaborative filtering. In: Proceedings of ACM SIGIR (2016)
16. Zhao, H., Yao, Q., Li, J., Song, Y., Lee, D.L.: Meta-graph based recommendation fusion over heterogeneous information networks. In: Proceedings of ACM SIGKDD (2017)
17. Zhao, X.W., Guo, Y., He, Y., Jiang, H., Wu, Y., Li, X.: We know what you want to buy: a demographic-based system for product recommendation on microblogs. In: Proceedings of ACM SIGKDD (2014)